

joinFS

By Matthew Harlan

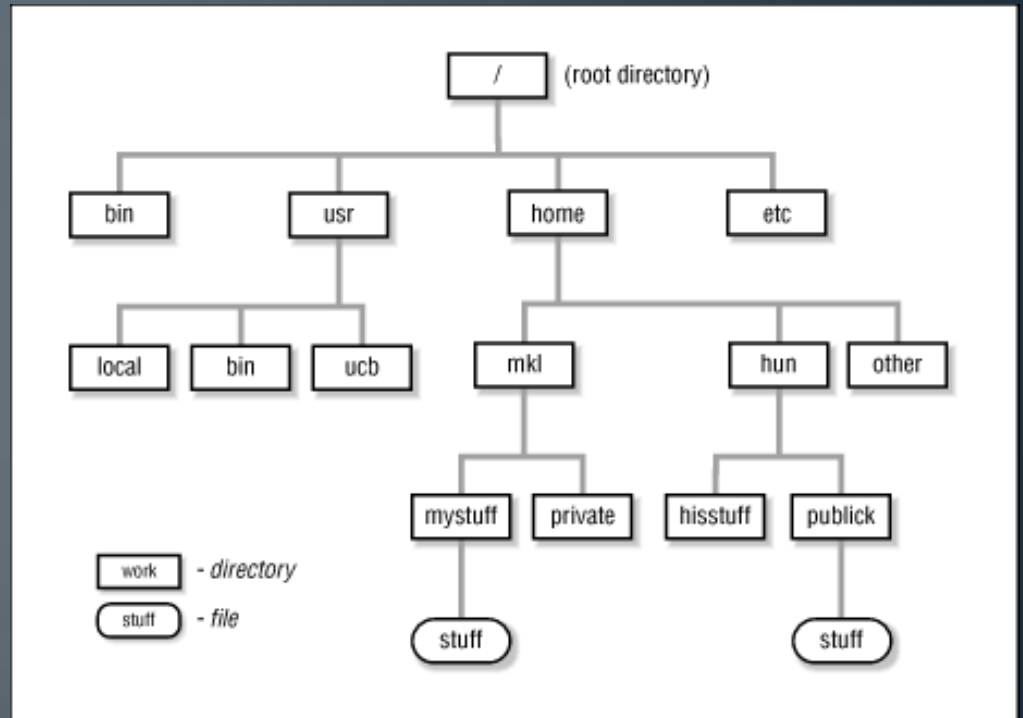
What is joinFS?

- A system for storing files on a computer
 - For Linux!
- A system for automatically organizing files

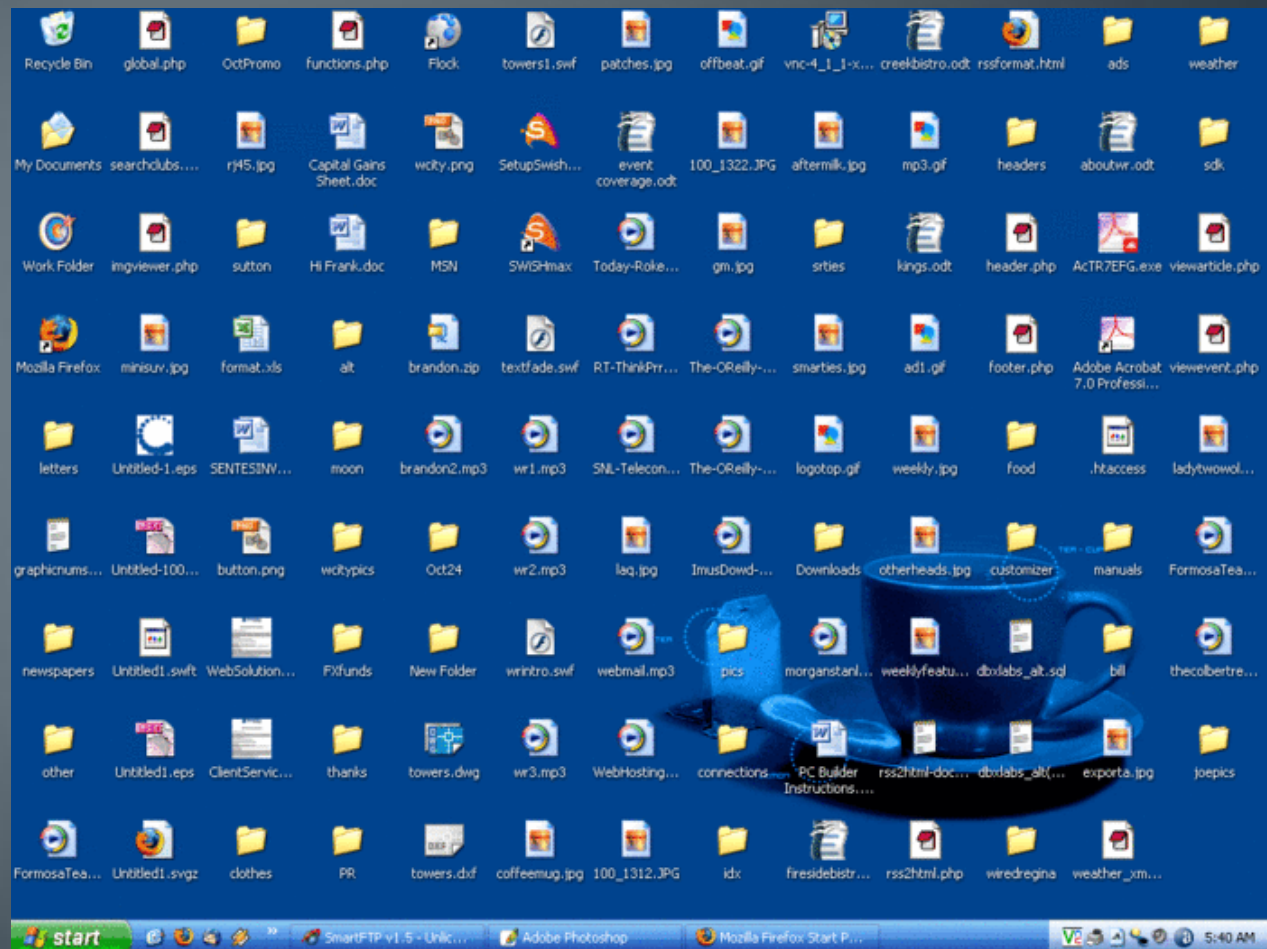


Traditional File Systems

- Folders categorize files
- Hierarchy of categories
- Managed by the user



Users Get Lazy




Hierarchical File Systems Are Dead

- Proclaimed in paper by Margo Seltzer and Nicholas Murphy
- Users have too much data to manage
- File systems must incorporate search



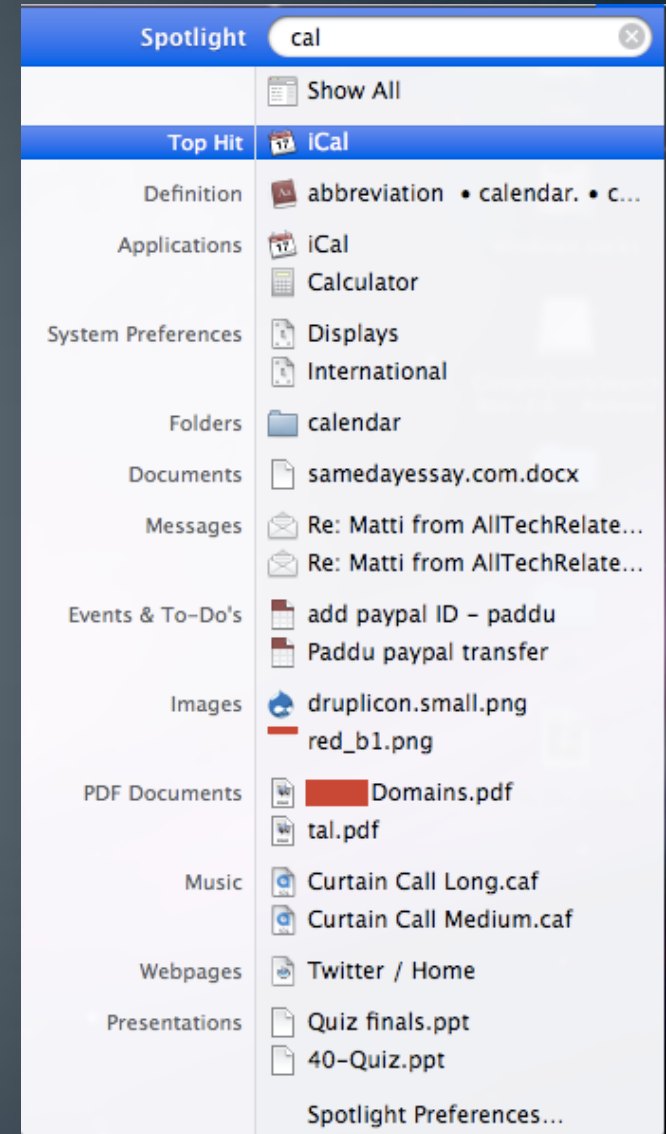
Files Are Complex Data

- Most files belong to multiple categories
 - Photos: time, place, people pictured, etc.



Current Solutions

- Windows Search, Spotlight
 - Additional program that runs on top of the file system
- Searches the file system and returns categorized results



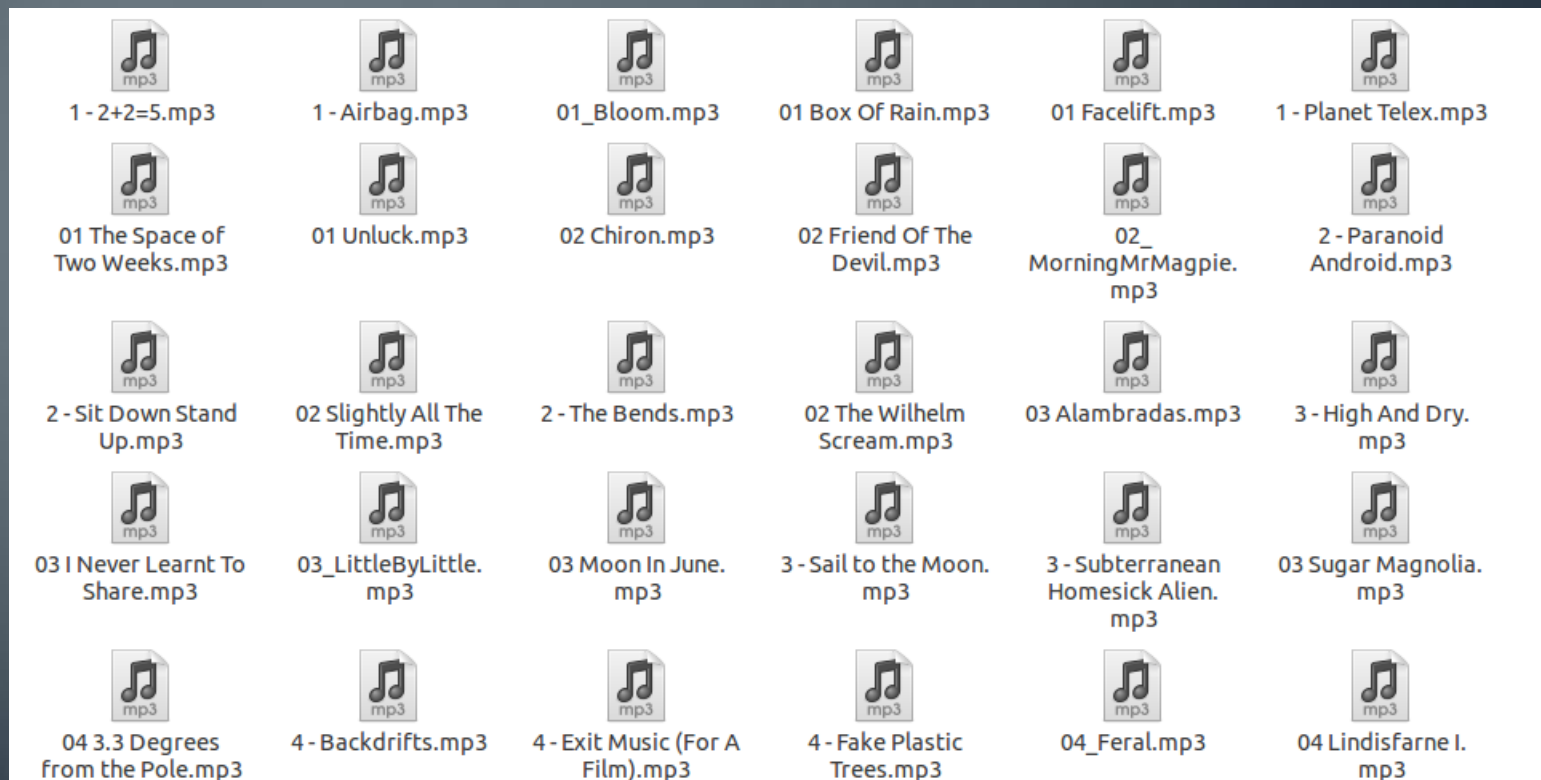
A Natural Next Step

- Represent searches as folders
 - Content is the result of the search!
- Enable users to search for any attribute
- Allow for a hierarchy of searches

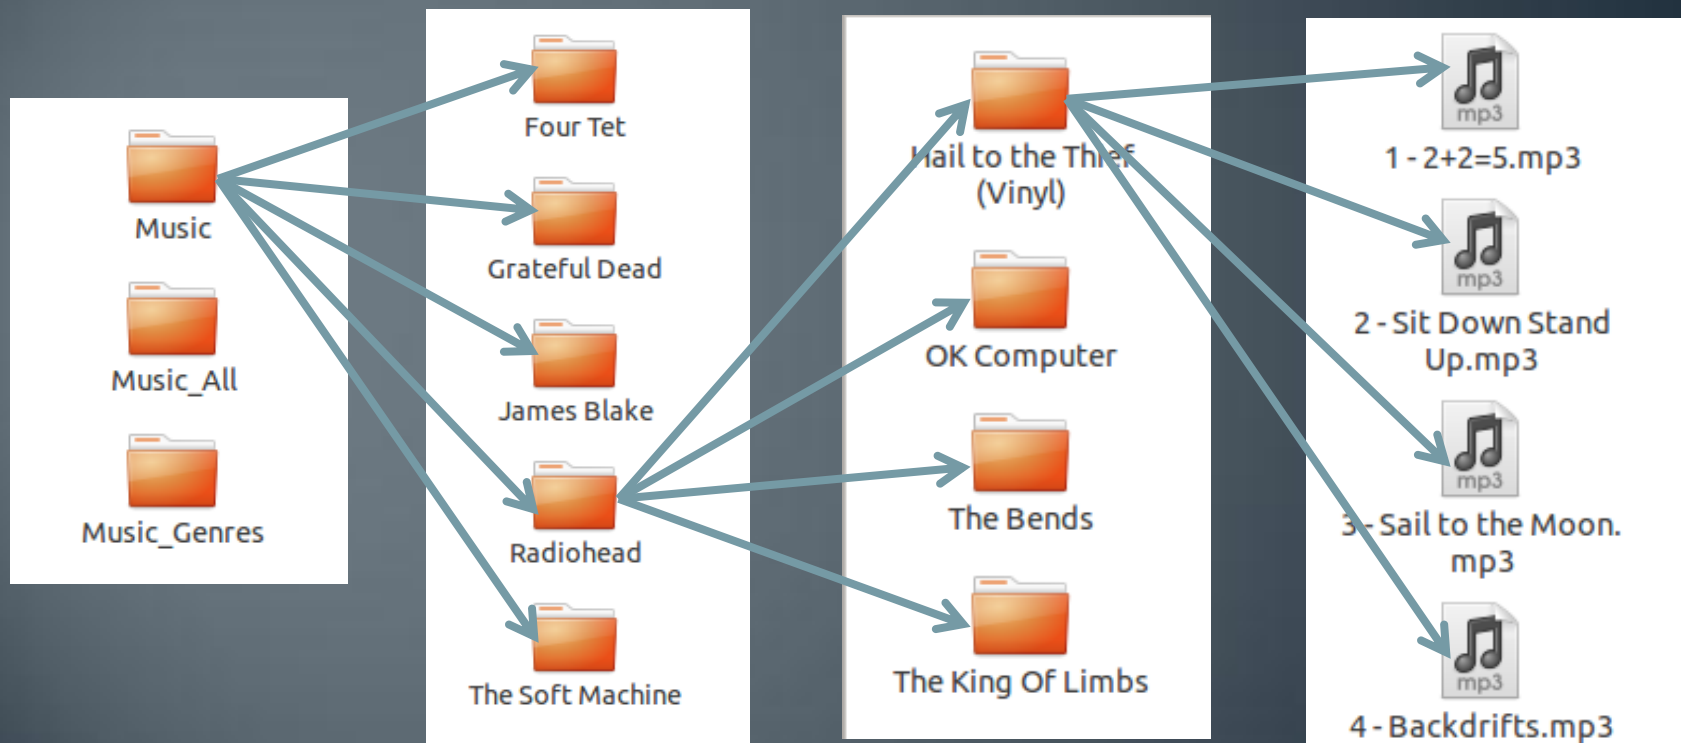


Search Folders

- Turns a folder of unorganized music files into an organized music library



Search Folders



The Demo

- A simple search for all music
- An automatically generated Music Library
 - Music > Artist > Album > MP3 files
 - Music Genres > Genre > Artist > MP3 files
- Delete and how it ripples through the hierarchy

Challenges


- Files and folders are stored on the underlying file system
 - Search results are not physically backed
- JoinFS must recognize a search result and return the file that it represents
- Must also maintain access restrictions

Backwards Compatibility

- **Files and Links:** getattr, readlink, mknod, unlink, symlink, rename, link, truncate, open, read, write, flush, release, fsync, create, ftruncate, fgetattr, lock, utimens
- **Extended Attributes (Meta Data):** setxattr, getxattr, listxattr, removexattr
- **Folders:** mkdir, rmdir, opendir, readdir, releasdir
- **Security:** chmod, chown, access
- **System:** init, destroy, stafs

Performance

- This is a real file system
 - Performance matters



- In February, joinFS took **40 minutes** to transfer about 40,000 files
- Now, it takes joinFS **30 seconds**

The Power of joinFS

- Uses pre-existing Linux file system commands
- Anything can be the result of a search
- You can use joinFS without understanding it
- Any program can use the search functionality



Any Questions?

- Source Code:
 - <http://www.github.com/mharlan/joinFS>